

A Pilgrimage through John Martin Fischer's *Deep Control: Essays on Free Will and Value*

I. Introduction

John Martin Fischer's most recent collection of essays, *Deep Control: Essays on Free Will and Value*, is incredibly wide-ranging and impressively detailed. Fischer manages to cover a staggering amount of ground in the free will debate, while also providing insightful and articulate analyses of many of the positions defended in the field. In this collection, Fischer focuses on the relationship between free will and moral responsibility. Traditionally, philosophers have taken there to be two requirements for moral responsibility—the control, or free will, condition and the epistemic condition.¹ In order for an individual to be morally responsible for a given act, that act must be under her control, or up to her in some way, and she must also know, or ought to have known,² that she is performing the act in question. In this collection of essays, Fischer focuses upon the control condition for moral responsibility. In the first section of his book, Fischer discusses perhaps the most widely held account of the control condition—regulative control. Regulative control requires access to alternative possibilities; in order to be responsible, agents must have the ability to do otherwise. While regulative control is a hugely influential account of the control condition, Fischer argues that certain cases, first described by Harry Frankfurt, show that the ability to do otherwise is irrelevant to our attributions of moral responsibility. Fischer spends the entirety of the first section of *Deep Control* defending Frankfurt cases against numerous attacks. Having thoroughly defended the claim that regulative control is unnecessary for moral responsibility, Fischer turns his attention to other accounts of control in the second section of the collection. Finding many compatibilist conditions for control lacking and many incompatibilist conditions too stringent, Fischer, taking

¹ Some philosophers also accept an autonomy or authenticity condition where, in order for an agent to be responsible for her behavior, she must in some way be the authentic source of her actions. For an account of such an authenticity condition, see Daniel Dennett's *Elbow Room: The Varieties of Free Will Worth Wanting* (1984). Many philosophers also argue that there is a tracing component to moral responsibility. By tracing back to a time when an agent fulfills the epistemic and control conditions, we can understand how an individual can be responsible for a given act even though she doesn't fulfill the epistemic or control conditions at the time she commits the act. Fischer, along with co-author Neal Tognazzini, tackles this issue in the essay, "The Triumph of Tracing" (Fischer 2012, 206-233).

² One could also defend a weaker epistemic condition, one that requires agents to justifiably believe that they are performing a certain act in order to be responsible for that act.

the “middle path,” defends his own account—guidance control. It really is quite astounding how rigorous a discussion Fischer provides in this collection and, while it is beyond the scope of this essay to discuss each chapter in full, every article offers insights into the nature of free will and moral responsibility. In this essay, I will first focus on Fischer’s defense of Frankfurt cases, specifically his response to the argument that the assumption of determinism in such cases is question-begging. Next, I’ll discuss Fischer’s account of guidance control and go on to analyze two objections to it. Finally, I will conclude with a brief discussion of the metaphor of the pilgrimage that Fischer introduces in the opening essay of this collection.

II. Frankfurt Cases and the Dilemma Defense

a. Fischer’s Frankfurt Case

Prior to Frankfurt’s essay “Alternate Possibilities and Moral Responsibility,” both incompatibilists and compatibilists alike largely accepted the principle of alternative possibilities (PAP):

PAP: a person is morally responsible for what she has done only if she could have done otherwise (Frankfurt 1969).

In this groundbreaking paper, Frankfurt constructs an example in which a counterfactual intervener eliminates an agent’s ability to do otherwise yet, intuitively, the agent remains morally responsible for what she does. Furthermore, the manner in which this counterfactual intervener renders the agent unable to do otherwise does not affect the agent’s process of deliberation, her beliefs and choices, or any other aspect of the agent. In this way, the agent acts on her own, uninfluenced by the intervener. This original Frankfurt case generated an enormous amount of literature, and has been subject to several modifications and revisions over the years. In this collection, Fischer presents his own Frankfurt-style example:

Because he dares to hope that Democrats finally have a good chance of winning the White House, the benevolent but elderly neurosurgeon, Black, has come out of retirement to participate in yet another philosophical example. (After all, what would these thought-experiments be without the venerable *éminence grise*—or

should it be *noir*?) He has secretly inserted a chip in Jones's brain that enables Black to monitor and control Jones's activities. Black can exercise this control through a sophisticated computer that he has programmed so that, among other things, it monitors Jones's voting behavior. If Jones were to show any inclination to vote for McCain (or, let us say, anyone other than Obama), then the computer, through the chip in Jones's brain, would intervene to assure that he actually decides to vote for Obama and does so vote. But if Jones decides on his own to vote for Obama (as Black, the old progressive, would prefer), the computer does nothing but continue to monitor—without affecting—the goings-on in Jones's head.

Now suppose that Jones decides to vote for Obama on his own, just as he would have if Black had not inserted the chip in his head. It seems to me, upon first thinking about this case, that Jones can be held morally responsible for his choice and act of voting for Obama, although he could not have chosen otherwise and he could not have done otherwise (Fischer 2012, 34).³

In the first section of his collection, Fischer emphasizes that Frankfurt cases are best utilized in a two-step argument. The first step is to argue that examples like these undermine the claim that moral responsibility requires the ability to do otherwise. Indeed, it certainly seems irrelevant to our judgment that Jones is morally responsible that the counterfactual intervener ensured the absence of alternative possibilities. So, if Jones fails to be responsible, it is not because the counterfactual intervener blocks his access to alternative possibilities. Once this is established, the compatibilist can then take the next step and argue that if the counterfactual intervener's elimination of alternative possibilities is irrelevant to moral responsibility, the fact that determinism also rules out access to alternative possibilities ought to be irrelevant as well. Thus, Fischer concludes, "if causal determinism rules out moral responsibility, it is not in virtue of eliminating alternative possibilities" (Fischer 2012, 36). Of course, even successfully completing this step of the argument is insufficient to defeat the incompatibilist; nothing about these cases rules out the possibility that determinism undermines some requirement for moral responsibility other than PAP. However, by rejecting PAP, Frankfurt defenders are able to widen the logical

³ The original version of this case can be found in Fischer (1982).

space in which compatibilists can stake their positions. It was previously thought that the dialectical burden of the compatibilist was to prove that determinism and PAP are compatible—a truly daunting task.⁴ However, with Frankfurt-style examples, the compatibilist is able to take an entirely different position—she can deny PAP.

b. The Dilemma Defense

Of course, not all are convinced that Frankfurt cases are able to rule out, or even call into question, PAP. Many argue that Frankfurt cases face a dilemma—either they beg the question against the incompatibilist or they fail to successfully expunge all alternative possibilities, and thus fail to undermine PAP.

This dilemma rests on the inclusion of a “prior sign,” a cue that lets the intervener know beforehand whether intervention is required.⁵ Prior signs play a very important role in Frankfurt cases. Without such a cue, it could be the case that the intervener, Black, would have to interfere *after* Jones had decided to vote for McCain (or anyone other than Obama) in order to ensure he votes for the right candidate. But if this is the case, then the defender of PAP can argue that these examples no longer pose a threat. After all, while it is true that Jones could not avoid voting for Obama, there would remain a morally significant respect in which Jones could have done otherwise—he could have chosen not to vote for Obama. Though necessary to get Frankfurt cases off the ground, the inclusion of these prior signs introduces a wrinkle for the Frankfurt defender. Many argue that the inclusion of the prior sign generates a dilemma: the prior sign is either an infallible predictor of Jones’s future behavior or it is not.⁶ If it is an infallible predictor, then the occurrence of the prior sign is sufficient for the occurrence of Jones’s future behavior. But this is to assume determinism, which begs the question against the incompatibilist who believes that determinism is incompatible with moral responsibility. But, if the prior sign does not infallibly predict Jones’s future behavior, then it looks as if he has alternative possibilities

⁴ Though many compatibilists maintain the two are compatible. See Lehrer (1976), Horgan (1979), Lewis (1981), Campbell (1997), and Vihvelin (2000).

⁵ Though Frankfurt does not include a prior sign in his original case, he does note the possibility of doing so in a footnote: “We can imagine that Jones₄ has often confronted the alternatives—A and B—that he now confronts, and that his face has invariably twitched when he was about to decide to do A and never when he was about to decide to do B. Knowing this, and observing the twitch, Black would have a basis for his prediction” (Frankfurt 1969, 835).

⁶ See Ginet (1996), Kane (1996), and Widerker (1995).

open to him. And, if the agent could do otherwise, then the case is not an adequate counterexample to PAP.

c. Fischer's Take on the Dilemma Defense

Many Frankfurt defenders choose to take on the indeterministic horn of the Dilemma Defense, leaving the deterministic horn largely abandoned.⁷ However, while Fischer has made incredibly important contributions to the literature on the indeterministic horn of the dilemma, he is unique in defending the assumption of determinism in Frankfurt cases. While a great deal of important and innovative work has been done to create cases in which alternative possibilities are unavailable without assuming the truth of determinism, such an avenue is not so easy to pursue. Regarding these modified Frankfurt cases, Michael McKenna notes: "I wish only to report that any example offered has been vigorously contested. It seems that there is no received opinion as to whether any of them will succeed... each involves intricacies that import further philosophically disputed theses" (McKenna 2008, 776). In light of these worries, Fischer's work on the deterministic horn of the dilemma becomes even more important—if it turns out that these modified indeterministic Frankfurt cases fail to be convincing, the compatibilist will need to be able to rely on the assumption of determinism in order to get Frankfurt cases off the ground, and they must be able to do so in a non-question-begging way.⁸ Luckily, Frankfurt defenders will be able to look to Fischer's work to do just this.

In earlier work, Fischer argues that his two-step argument based on deterministic Frankfurt cases can be presented in a non-question-begging way. As stated earlier, the first conclusion to draw when presented with a deterministic Frankfurt case is that *if* the agent isn't morally responsible, it's not because he doesn't have access to alternative possibilities (Fischer 1999; 2004). From here, the Frankfurt defender can argue that if the counterfactual intervener's elimination of alternative possibilities fails to undermine our ascription of moral responsibility, then the truth of determinism should fail in exactly the same respect. This conclusion doesn't beg

⁷ For examples of those who take on the second horn of the Dilemma Defense, see Hunt (2000), McKenna (2003), Mele and Robb (1998), and Pereboom (2000).

⁸ Furthermore, there are some compatibilists who contend that we cannot be morally responsible unless determinism is true or, less demandingly, unless deterministic causal relations can obtain. For example: Ayer (1954), Hobart (1934), Nowell-Smith (1948), and Smart (1961). These theorists in particular must be able to respond to the first horn of the Dilemma Defense.

the question against incompatibilists because it could still be the case that the truth of determinism rules out moral responsibility—it would simply do so in a way that has nothing to do with ruling out alternative possibilities.

However, many have challenged this initial conclusion. After all, if determinism is true in Fischer's Frankfurt case, then, according to some incompatibilists, Jones never had access to alternative possibilities, regardless of Black or Black's device. Whatever prior sign Jones exhibits to indicate that he will choose to vote for Obama—a raised eyebrow or a slight nose twitch—is sufficient for Jones to choose to vote for Obama in a deterministic world. Neither Black nor Black's machine ruled out the possibility that Jones votes for someone other than Obama—such a possibility was ruled out as soon as Jones's eyebrow rose (and, of course, as these critics see it, long before that, really). And if this is the case, then the Frankfurt defender hasn't shown that Jones's lack of alternative possibilities is irrelevant to whether he is morally responsible for voting for Obama. Rather, she has simply stated, rather flat-footedly, that even though determinism rules out access to alternative possibilities, this is irrelevant to moral responsibility, which is, of course, precisely what the incompatibilist and defender of PAP wish to deny.

In Chapter 2, "The Frankfurt Cases: The Moral of the Stories," Fischer accepts this criticism, and modifies the presentation of his deterministic Frankfurt case to address this worry. First, he argues that one can present a deterministic Frankfurt case while remaining agnostic about whether determinism rules out alternative possibilities. In so doing, the Frankfurt defender can render counterfactual interveners efficacious once again. In Fischer's example, once Black sees the prior sign that indicates that Jones will vote for Obama, he can back off, for he knows that, because determinism obtains, the prior sign is sufficient for Jones to choose to vote for Obama. If Jones failed to exhibit the relevant prior sign, then Black would intervene. According to Fischer, in this presentation of the case, both the assumption of determinism *and* Black's ability to act on whichever prior sign Jones exhibits make it the case that Jones couldn't do otherwise than vote for Obama. Below is the formal presentation of Fischer's argument based on this modified Frankfurt case:

1. Assume that causal determinism obtains and that the Frankfurt case of Jones and Black unfolds as above.

2. At this point in the argument, causal determinism is not assumed in itself to rule out access to alternative possibilities. (Neither is it supposed that Black's presence, device, and dispositions in themselves rule out such access.)
3. Causal determinism plus Black's presence, device, and dispositions rule out Jones's freedom at t2 to choose otherwise.
4. If Jones is not morally responsible for choosing at t2 to vote for Obama at t3, it is not in virtue of the *mere* fact that he was not free at t2 to choose otherwise.
5. If causal determinism rules out Jones's moral responsibility for his choice at t2, it is *not* in virtue of its eliminating alternative possibilities (if in fact it does eliminate alternative possibilities) (Fischer 2012, 44).

As Fischer sees it, by remaining agnostic about whether determinism rules out access to alternative possibilities, he has endowed Black with the ability to constrain Jones's options. And, if we can agree that the counterfactual nature of Black's intervention is irrelevant to whether Jones is morally responsible, then the compatibilist can make the further argument that access to alternative possibilities is irrelevant to our ascriptions of moral responsibility.

d. A Challenge to the Agnostic Assumption

Fischer's modified Frankfurt case is creative and poses a real challenge to those who dismiss deterministic Frankfurt cases as question-begging and unconvincing. The key move in this argument is making the agnostic assumption—remaining neutral about whether determinism eliminates alternative possibilities. Without this assumption, Dilemma Defenders are able to argue that counterfactual interveners, such as Black, are merely inefficacious red herrings, while determinism alone eliminates alternative possibilities. In this section, I argue that this assumption, while vital to the success of this argument, cannot be stably held. Though the agnostic assumption may make room for Black at the table as a source of alternative possibility elimination, it also places him in a precarious position, one that is ultimately detrimental to the Frankfurt defender.

When explaining how Black, along with the truth of determinism, are able to rule out alternative possibilities, Fischer argues: “Given that Black knows that causal determinism obtains, he can now relax, as it were; under these circumstances, Black knows that Jones in fact will subsequently choose to vote for Obama and carry out that choice” (Fischer 2012, 42). But if we are operating under the agnostic assumption, how could Black *know* that Jones will vote for Obama once he witnesses the prior sign? Because we are remaining agnostic about whether determinism rules out alternative possibilities, neither we, nor Black, can conclude that Jones cannot do anything other than vote for Obama once the prior sign is exhibited. It seems as though the only way to stably hold the agnostic assumption is to grant that alternative possibilities could be open to Jones even after he exhibits the prior sign. But this places Black in a precarious position. If alternative possibilities could be open to Jones after the prior sign, but Black doesn’t intervene, then the Dilemma Defender can argue that Jones very well may have had access to alternative possibilities and *this* is why we judge him to be morally responsible for voting for Obama. But, if Black does intervene, though he would have certainly taken part in eliminating alternative possibilities, any actual intervention would undermine our ascriptions of moral responsibility. And, if the Frankfurt defender simply denies that alternative possibilities could exist after the prior sign is exhibited, then it looks as though the agnostic assumption has been jettisoned and Black is again relegated to the role of a red herring, for the prior sign is what necessitates Jones’s behavior. In this way, the agnostic assumption is both a blessing and a curse. It’s a blessing in that it gives the counterfactual intervener, Black, some work to do. But it’s a curse in that the mere presence of Black is not enough to do that work—Black must actually intervene to eliminate alternative possibilities when we operate under the agnostic assumption. And, of course, actual intervention would undermine the goal of Frankfurt cases, for it generates the intuition that Jones is not morally responsible.

Of course, one can easily make room for Black to rule out alternative possibilities in Frankfurt cases by jettisoning the assumption of determinism. Indeed, all of the key features of Fischer’s modified Frankfurt case can be preserved if one assumes indeterminism.⁹ And while this would be an admirable response to the indeterministic horn of the Dilemma Defense, it leaves the deterministic horn unanswered. And, as Fischer himself argues, it’s crucial for

⁹ Fischer argues that this case would be similar to David Hunt and Derk Pereboom’s “Buffer Zone” cases, which are presented in response to the indeterministic horn of the Dilemma Defense (Hunt 2000; Pereboom 2001).

compatibilists to be able to defend the assumption of determinism in Frankfurt cases. Without a proper defense of the assumption of determinism, it becomes significantly harder for compatibilists to convince undecided parties of the irrelevance of PAP. Though it may be the case that the agnostic assumption is not the best way to defend deterministic Frankfurt cases, Fischer's work in this area offers great insight into both the importance and difficulties of defending these cases.

e. Other Objections to Frankfurt Cases

Though the Dilemma Defense is perhaps the most common strategy in objecting to Frankfurt cases, Fischer provides responses to several other objections, presenting a truly developed defense of these cases. While it is beyond the scope of this review to provide full analyses of these objections and responses, I will provide brief synopses of each, if to do nothing more than illustrate how very thorough Fischer's collection truly is.

In Chapter 3, "Freedom, Foreknowledge, and Frankfurt: A Reply to Vihvelin," Fischer responds to Kadri Vihvelin's recent challenges to the logic of Frankfurt cases. Vihvelin first argues that agents, such as Jones, in Frankfurt cases never lose the ability to do otherwise, even given the presence of a counterfactual intervener. To illustrate this claim, Vihvelin introduces a case, presented in an abbreviated form below:

An agent, let's call him Jones, bets another individual that a coin will come up heads. It does so and Jones wins the bet. In this case there is also a confederate, Black, who is capable of predicting the outcomes of coin flips with incredible accuracy, predicts that the coin will come up heads several hours before the coin flip takes place. Black is friends with Jones and wants him to win, so if he had predicted that the coin would come up tails, he would have interfered in some way to make it nomologically necessary that the coin would come up heads. But, in this case, he does nothing.¹⁰

According to Vihvelin, Jones wins the bet fair and square. But, importantly, even though the coin will never come up tails, that doesn't mean that it *cannot*. Vihvelin then argues that this case is analogous to Frankfurt cases. In Fischer's Frankfurt case, for example, though Jones never *will*

¹⁰ Vihvelin's full case can be found in *Deep Control* on p. 56 and in Vihvelin (2000).

choose to vote for someone other than Obama, this doesn't mean that he *can't* choose otherwise. Fischer responds by arguing that there is a relevant difference between the coin case and his Frankfurt case. In the coin case, because Black is far away when the coin flip occurs, if the coin were about to come up tails, it would come up tails.¹¹ But in the Frankfurt case, if Jones were to exhibit the prior sign that indicates he is about to choose to vote for someone other than Obama, Black *will* intervene. In other words, even if Jones is about to choose to vote for someone other than Obama, he will be unable to. Yet, in the coin case, if the coin is about to come up tails, it will. So, even if we grant that the coin really could come up tails in Vihvelin's coin case, we can still maintain that Jones cannot do otherwise in the Frankfurt case.

However, Vihvelin argues that the counterfactual, on which Fischer relies to distinguish between the coin case and his Frankfurt case, is logically suspect. Below is Vihvelin's reconstruction of Fischer's argument:

If [the relevant agent] were about to refrain (in the absence of intervention by an external agent or factor), the triggering event would already have occurred.

If the triggering event had already occurred, Black would have intervened and forced Jones to act, in which case Jones would not have been able to refrain.

Therefore, if Jones were about to refrain, he would be rendered unable to refrain (Fischer 2012, 63).

According to Vihvelin, this counterfactual reasoning is a hypothetical syllogism, which is an invalid form of argument. For example, the following hypothetical syllogism is clearly invalid: "If I jumped off this bridge, then I would have arranged to wear a parachute. If I were wearing a parachute, I would not be killed. So if I jump off this bridge, I would not be killed" (Vihvelin, 20). However, Fischer argues, even in Vihvelin's presentation of the case, there is no problematic "world-hopping" that is characteristic of other invalid uses of hypothetical syllogism. The world in which the conclusion is true is the very same world in which the

¹¹ If this were to occur, then Black would have a false belief about the coin flip. Fischer grants that Black can have knowledge about the future, but this does not mean that Black is infallible. If he were incapable of having false beliefs, then this would call into question whether the coin really could have come up tails. Rather, Black has human foreknowledge, as opposed to divine foreknowledge. He can know contingent propositions; this doesn't make these propositions necessary. Fischer clarifies: "Although the *conditional* 'If Black knows that *p*, then *p* is true' can be necessary, this necessity does not attach to the consequent of the conditional, even given the truth of the antecedent" (Fischer 2012, 59).

premises are true (Fischer 2012, 64). Furthermore, Fischer argues that even if the form of hypothetical syllogism alone cannot provide a valid argument, other facts about the Frankfurt case license the inference to the conclusion (Fischer 2012, 61). Remember that Black is completely reliable in predicting what Jones will be about to do in the future, and will intervene if the prior sign indicates that Jones is about to choose to vote for someone other than Obama in the future. Given these facts, it just seems obvious that Jones is unable to do otherwise. To get to this conclusion, we don't need to generate an argument using hypothetical syllogism. Fischer concludes that there is no threat to the logic of Frankfurt cases.

In Chapter 4, "The Importance of Frankfurt-style Argument," Fischer responds to an argument by Daniel Speak. Speak, in his essay, "The Impertinence of Frankfurt-style Argument" (2007), concedes that Frankfurt cases really do undermine PAP, but argues that they are completely irrelevant to the debate between compatibilists and incompatibilists. In Frankfurt's original analysis of these cases, he argues that agents like Jones are morally responsible because even though they couldn't in fact do otherwise, if they were able to do otherwise, they would have behaved exactly as they actually did. Speak argues that this requirement of counterfactual stability will not do, however, for the truth of determinism rules out such counterfactuals. In order to determine what the agent would do if she could have done otherwise, we must locate a possible world in which such conditions obtain. But the closest possible worlds are indeterministic, and it's unclear how an agent would behave in them. So, we really can't know that an agent would perform the same action that she does in this world even if she could do otherwise. Thus, Jones is not morally responsible in Fischer's Frankfurt case, for determinism rules out counterfactual stability, which even Frankfurt thought was required for moral responsibility. In response, Fischer argues that, while Frankfurt did indeed provide such counterfactual analysis, he needn't, and shouldn't, have done so. Indeed, Frankfurt himself offered at least two other explanations for why agents like Jones are morally responsible, neither of which relies on counterfactual stability.

Nevertheless, Speak continues to argue that Frankfurt cases, though they successfully undermine PAP, do not pose a threat to incompatibilists. To do this, Frankfurt cases must illustrate the falsity of PAP *in the right way*. But, according to Speak, Frankfurt cases cannot do this, for while counterfactual interveners play no role in the behavior of the relevant agents, determinism does: it exerts causal influence over agents' behavior. In response, Fischer

articulates his two-step strategy for deploying Frankfurt cases. The first step, which Speak grants, illustrates the falsity of PAP—the ability to do otherwise is irrelevant to moral responsibility. Now, if alternative possibilities are irrelevant to moral responsibility, then Fischer makes the next step and asks: Why would it matter *how* these alternative possibilities are eliminated if Speak grants that their absence is irrelevant to moral responsibility? If one grants that the disappearance of alternative possibilities is irrelevant, how they disappear must also be irrelevant. Thus, Fischer concludes that Frankfurt cases undermine PAP *and* show that determinism doesn't rule out moral responsibility by eliminating access to alternative possibilities. Of course, Fischer is eager to grant that Frankfurt cases are not enough to fully undermine the incompatibilist's position—after all, determinism could rule out moral responsibility in a way that has nothing to do with alternative possibilities. But Frankfurt cases are still important, for they are able to undermine the most prominent *indirect* argument against compatibilism.

In Chapter 5, Fischer and his co-author Neal Tognazzini respond to an essay by Michael Otsuka in which he argues that even if Frankfurt cases undermine PAP, they do not undermine a related principle: the “Principle of Avoidable Blame (PAB).”

PAB: One is blameworthy for performing an act of a given type only if one could instead have behaved in a manner for which one would have been entirely blameless (Otsuka 1998).

According to Otsuka, this principle entails incompatibilism and is not threatened by Frankfurt cases. While PAB can be interpreted in many ways, Fischer and Tognazzini argue that Otsuka most likely intended to defend a strong reading of the principle, one on which an agent must be able to voluntarily perform both the actual blameworthy act and the counterfactual blameless act. However, on this interpretation, PAB falls victim to Frankfurt cases. After all, the agent could never voluntarily pursue the counterfactual blameless action, for the counterfactual intervener would have already eliminated this as a possible alternative, long before the agent could have voluntarily chosen to pursue it. And yet it really does seem as though agents in such cases are morally responsible for their behavior.

Though Otsuka is skeptical that a counterfactual intervener really could close off the possibility of acting blamelessly without influencing the agent in question, he argues that even if

such a case could be constructed, the agent would not be blameworthy. According to Otsuka, in order to be blameworthy, an agent either must have known or should have known that he or she was able to behave less badly than he or she actually does. But because knowledge requires true belief, it would be impossible for agents in Frankfurt cases to know that they could have behaved less badly, for it was impossible for them to do so. Thus, Otsuka argues that these Frankfurt agents are not blameworthy for their behavior. In response, Fischer and Tognazzini argue that there is no reason to accept Otsuka's knowledge condition for blameworthiness. After all, there are a variety of competing compatibilist conditions for blameworthiness that are met in Frankfurt cases. Thus, in order to argue convincingly that Frankfurt agents are not blameworthy, Otsuka must be able to object to the blameworthiness conditions that are met in these cases. Fischer and Tognazzini conclude that PAB fails to be more successful in the face of Frankfurt cases than PAP.

In Chapter 6, "Indeterminism and Control: An Approach to the Problem of Luck," Fischer turns his attention to William James's "Dilemma of Determinism." Fischer reconstructs the dilemma in the following way:

1. Either causal determinism is true, or it is false.
- 2*. If causal determinism is true, then I cannot do otherwise, and thus I am not morally responsible for my actions.
- 3*. If causal determinism is false (in a relevant way, i.e., in the sequences leading to my behavior), then my actions are not appropriately connected to my prior states (i.e., "my actions" are not in a genuine sense *my* actions), and thus I am not morally responsible for my actions.

Therefore:

4. I am not morally responsible for my actions (Fischer 2012, 85-86).

To respond to the deterministic horn, Fischer grants that determinism rules out the ability to do otherwise, but argues that the ability to do otherwise is unnecessary for moral responsibility. To illustrate this claim, he returns to his now familiar Frankfurt case. Rather than focusing on the deterministic horn of the dilemma, Fischer turns his attention to the indeterministic horn and the problem of luck.

The most prominent argument for the third premise is the “Rollback Argument,” which calls into question whether agents can have the right relationship to their actions in a deterministic world. In order for an individual to be morally responsible for an action, it must be *hers*—it must be related to her prior mental states in the appropriate way (Fischer 2012, 91). But if indeterminism is true, then an agent’s actions are simply matters of luck—they are not appropriately related to her mental states (her desires, beliefs, intentions, etc.). If we were to somehow rollback the universe to a time before the agent made a certain decision, it is unclear that she would make that same decision again. Thus, indeterminism undermines the responsibility-grounding relation necessary for our ascriptions of moral responsibility.

In response to the Rollback Argument, Fischer questions whether indeterminism really is capable of undermining the responsibility-grounding relation. He asks the reader to imagine two worlds. In the first world, W1, Jones chooses to raise his hand at t_2 and actually raises his hand at t_3 . Let us suppose that the responsibility-grounding relation obtains in this world. In the second world, W2, everything is exactly as it is in W1, except that there exists a randomizing machine, which could, at time $t_{1.9}$, either initiate a causal sequence that would stimulate Jones’s brain in such a way that he would not raise his hand at t_3 or do nothing. In fact, the machine does not interfere, so Jones chooses to raise his hand at t_2 and does so at t_3 , just like Jones does in W1. Fischer argues that if the responsibility-grounding relation obtains in W1, then it must also obtain in W2. Surely the mere *presence* of this machine cannot undermine the responsibility-relation; the machine has no effect on the causal sequence in W2, which is identical to the sequence in W1. In W1, the responsibility-grounding relation obtains and, according to Fischer, it does so because of the relevant causal sequence. Indeed, Fischer argues that whatever grounds our ascriptions of responsibility must be a feature of actual causal sequences. Thus, if the causal sequence is identical in W2, the responsibility-relation ought also to obtain. The random machine is merely an “untriggered preemptor” and as such cannot affect the actual causal sequence, the very thing that Fischer thinks determines our attributions of responsibility. Thus, Fischer concludes that the Rollback Argument cannot support premise 3 of the determinism dilemma—indeterminism does not entail that actions cannot be relevantly connected with agents.¹² Fischer

¹² Fischer also responds to arguments for premise 3 of the dilemma defended by Ishtiyaque Haji, Robert Kane, and Alfred Mele (Fischer 2012, 99-101). Though it is beyond the scope of this review essay to discuss the specifics of these objections and replies, Fischer provides an excellent and comprehensive discussion of the problem of luck.

concludes this chapter by re-enforcing the claim that, when it comes to moral responsibility, we must focus on the actual causal sequences that lead to people's choices and actions; the existence of counterfactual interveners and untriggered preemptors ought not affect our judgments of moral responsibility.

The adeptness with which Fischer is able to articulate these objections and subsequently dismiss them is truly impressive. Yet Fischer is not satisfied with completing only this task. Once he is able to conclude that the ability to do otherwise is unnecessary for moral responsibility, Fischer focuses on positive accounts of the control condition, including his own.

III. Guidance Control

In the first half of his collection, Fischer argues that "regulative control," which requires genuine metaphysical access to alternative possibilities, is irrelevant to our attributions of moral responsibility. In the second half of his collection, Fischer turns his attention to accounts of what *is* required to be in control in order to be morally responsible. First, he responds to several objections to compatibilist accounts of moral responsibility. In Chapter 7, "The Direct Argument: You Say Goodbye, I Say Hello," Fischer responds to various versions of the direct argument, concluding that many versions rely on problematic modal principles and those that do not rely on such principles are wholly unappealing. In Chapter 10, "Sourcehood: Playing the Cards That Are Dealt You," Fischer focuses on arguments contending that moral responsibility requires agents to be the ultimate source of their actions and behavior. Fischer argues that such requirements are much too demanding; we can be morally responsible for our actions, as well as the source of our actions, without being self-created. Fischer also finds several compatibilist accounts of the control condition wanting. In Chapters 8, "Conditional Freedom and the Normative Approach to Moral Responsibility," and 9, "Judgment-Sensitivity and the Value of Freedom," he argues that conditional accounts of control and those that rely on judgment sensitivity are too shallow, for both accounts are susceptible to Frankfurt-style arguments. When positing his own account, Fischer seeks the middle ground between the impossibly demanding incompatibilist conditions for control and the freedom-irrelevant compatibilist conditions. Fischer argues that "guidance control" is what we care about when it comes to moral responsibility. In this section, Fischer does a thorough job of clearing the ground for his own

theory, and while this work is illuminating, I will only focus on Fischer's positive account in this essay. First, I will briefly introduce the theory of guidance control and then go on to analyze two objections to it.

a. Fischer's Account of Guidance Control

According to Fischer, in order for an agent to exercise guidance control over a particular action, the mechanism that issues that action must be moderately reasons-responsive.¹³ Moderate reasons-responsiveness involves being normally reasons-receptive (able to recognize reasons as such) and at least weakly reasons-reactive. According to Fischer, an agent is weakly reasons-reactive if the mechanism that issues her actual behavior would issue different behavior in a possible world in which the agent had sufficient reason to behave differently (Fischer 2012, 168). Notice that though we must undergo a counterfactual analysis to determine reasons-reactivity, it is only to determine if decision mechanisms possess certain modal properties. Importantly, there is no need to determine whether an agent is able to behave differently in any given situation. For Fischer, moral responsibility is tied to actual sequences of events.

In addition to being reasons-responsive, the relevant mechanism must also belong to the agent in order for her to exercise guidance control. After all, an agent could be reasons-responsive because she was brainwashed, manipulated, or tampered with in some other way. So, Fischer argues, an agent's mechanism must also be her own. Fischer defends a subjective approach to ownership—an agent owns a mechanism by taking the actions that stem from it as belonging to her. An agent must identify with her choices and actions; she must consider them hers. Thus, according to Fischer, an agent exercises guidance control, and is thus morally responsible, if she is moderately reasons-responsive and takes ownership over her mechanisms.

b. Mele's Agoraphobe

Fischer's account of guidance control has generated a great deal of discussion in the literature, and with good reason. Alfred Mele, in his essay "Reactive Attitudes, Reactivity, and Omissions"

¹³ Where a "mechanism" is to be understood as a process or way of making decisions and producing actions (Fischer 2012, 186).

(2000), presents several objections to Fischer's account of moderate reasons-responsiveness. Mele argues that moderate reasons-responsiveness may be too weak a condition for guidance control and moral responsibility. For example, those who suffer from extreme phobias may be reasons-receptive and weakly reasons-reactive, but we wouldn't want to hold them morally responsible for their actions. Mele asks his readers to imagine an extreme agoraphobe, who chooses to stay home rather than attend his daughter's wedding at a nearby church (Mele 2000, 288). Though he stays home in the actual world, and many close-by possible worlds as well, there is at least one possible world in which the agoraphobe's house catches on fire, and because he is even more afraid of fire than of leaving his home, he decides that he has good reason to leave, and does so. Because there is at least one world in which the agoraphobe would behave differently, given sufficient reason to do so, he is weakly reasons-reactive. And given that he feels ownership over his behavior, the agoraphobe meets all of the conditions necessary for guidance control. However, Mele argues that the agoraphobe is intuitively not morally responsible for missing his daughter's wedding.

Fischer accepts the criticism and argues that his account of weak reasons-reactivity should be developed such that those at the outer reaches of reasons-reactivity would not be considered morally responsible. Fischer grants that the agoraphobe is weakly reasons-reactive, but argues that he is not reasons-reactive enough to be able to exercise guidance control. Though Fischer does not elaborate on exactly how reasons-reactive one must be to cross such a threshold, and it may be very difficult to develop a precise account of minimally sufficient reasons-reactiveness, it seems largely uncontroversial that an agent can be so weakly reasons-reactive that he fails to be an appropriate target for our attributions of moral responsibility.

However, I wonder exactly what feature of the agoraphobe's case rules him out of the reasons-responsive community. Indeed, at no point does Fischer explicitly state what it is about the agoraphobe that makes him an inappropriate target of our responsibility attributions. Is it because there are only a few worlds in which he would behave differently, given different reasons? Or does it have to do with the nature of the reasons for which he behaves? Is there something particularly responsibility-undermining about overwhelming fears, for example? Though Fischer focuses more intently on how agents react to the reasons they have, there are some circumstances in which the reasons themselves affect a mechanism's ability to be reasons-

reactive. There's a question about which plays the more important role in reasons-reactivity—reasons or reactivity. The case discussed in the next section brings this point into focus.

c. Long's Tie-Breaking Voter

Todd Long, in his essay, "Moderate Reasons-Responsiveness, Moral Responsibility, and Manipulation" (2004), presents a challenge to guidance control in the form of a manipulation case. Long describes a man named Schmidt who is given the task of casting the tie-breaking vote in the debate over whether Hitler should take over Germany. Schmidt is conflicted—Hitler and he began their careers in politics around the same time, and Schmidt had great respect for Hitler, impressed by his early plans to create a better Germany. However, Schmidt began to hear disturbing stories about Hitler's plans and motivations in later years. As he's going to cast his vote, Block, who is very good at predicting what people will do, decides that he'll intervene if he thinks that Schmidt will not vote for Hitler. If Schmidt decides on his own to vote for Hitler and then does so, it seems clear that he is both responsible and blameworthy for his behavior, and Fischer's account of guidance control can easily accommodate this intuition. But if Schmidt were going to vote against Hitler, and Block had to intervene, it seems as though Schmidt ought not be held responsible for his behavior. Long argues that this intuition will prove troublesome to explain for Fischer.

Of course, there are many ways Block could manipulate Schmidt. Block could, for example, tamper directly with Schmidt's central nervous system, detouring his decision mechanisms. If this were the case, Fischer can accommodate the intuition that Schmidt isn't morally responsible; the mechanism that issued the choice to vote for Hitler was neither reasons-responsive nor Schmidt's own. However, Long presents another manipulation scenario in which Fischer's conditions for guidance control seem to be met:

A minute before Schmidt casts his vote, Block becomes convinced that Schmidt is going to vote against Hitler. Block goes into action: he *adds* new inputs into the *very same mechanism* that is operant in the actual sequence. Suppose that these new inputs come in the form of *reasons* for voting for Hitler. Block directly feeds into Schmidt's deliberative voting mechanism enough reasons, or reasons powerful enough, to ensure that Schmidt will vote for Hitler (Long, 158).

Long also presents a third scenario in which Block removes inputs from Schmidt's voting mechanism, rendering the weight of Schmidt's reasons heavily in favor of voting for Hitler. In these cases, it looks as though Fischer is committed to the claim that Schmidt is morally responsible. After all, merely adding or subtracting reasons from a mechanism doesn't change the mechanism, and if the mechanism was reasons-responsive in the actual sequence, it would be reasons-responsive in these scenarios as well. Interestingly, Long doesn't think these scenarios are counter-examples to Fischer and Ravizza's account. Rather, he thinks that Schmidt really is morally responsible in these scenarios, despite one's first reaction. According to Long, we ought to hold Schmidt responsible for voting for Hitler even if Block had either added reasons to, or subtracted them from, Schmidt's decision mechanism. After all, Schmidt was voting for the reasons he *had*, which is the responsible thing to do. While one could argue that the way Schmidt became aware of these reasons is suspect, it would surely be too heavy a burden to expect agents to be responsible not only for their choices and actions, but also for the reasons for which they act and how they acquire such reasons. Long argues:

Evidence (and reasons) supporting beliefs come to us from various directions and in many different ways. Some of it we are quite aware of, and some of it we are less aware of... Schmidt could have been convinced to vote by more ordinary external means (news reports, personal testimony, etc.). What is the relevant difference between manipulating the inputs in those ways and manipulating the inputs directly (Long, 165)?

According to Long, though responsibility is a matter of acting for reasons, the *source* of those reasons is immaterial to our responsibility judgments.

Fischer is unwilling to accept this explanation. Rather, he argues that Schmidt does not exercise guidance control in Long's manipulation cases. According to Fischer, Schmidt is not responsible in these cases because the mechanism that produces his choice to vote for Hitler is not reasons-responsive—it's an altogether different mechanism than the one utilized when Schmidt decided on his own to vote for Hitler. While Long is dubious that a mechanism can change simply by adding or taking away inputs, Fischer contends that this is exactly what occurs in the manipulation cases. He argues:

When ‘inputs’ are implanted in a way that does not allow for a reasonable or fair opportunity to subject those inputs to critical scrutiny in light of the agent’s normative orientation, then such manipulation does indeed remove moral responsibility. Crucially, I claim that such manipulation typically ‘changes the mechanism’ (Fischer 2012, 197).

According to Fischer, if Block implanted the reasons to vote for Hitler in such a way that Schmidt’s voting mechanism had the power to evaluate them, then Schmidt may very well be morally responsible for his subsequent behavior. But, as Long’s cases seem to suggest, if Block implanted the reasons only moments before Schmidt was due to vote, he couldn’t possibly weigh this information properly in such a short period of time.¹⁴ While Fischer agrees with Long that it is too demanding to require agents to be responsible for how they acquire reasons, once they are acquired, agents’ mechanisms must be able to filter them in order to be properly reasons-responsive.

d. *Reasons-Responsiveness vs. Reasons-Responsiveness*

While Fischer is clear that Schmidt’s voting mechanism isn’t reasons-responsive in the manipulation cases, it’s unclear why. At points, Fischer seems to be arguing that there simply isn’t enough time in the manipulation cases for Block’s inputs to be properly weighed. And though Fischer is most surely right that a mechanism could never be considered reasons-responsive if it is unable to filter reasons properly, it’s unclear that such a process would take much time, if any. Indeed, it’s an important feature of human agency that we don’t need to spend much time evaluating reasons to know that they are reasons for which we should act. Once a person, like Mele’s agoraphobe, realizes his house is on fire, it doesn’t (or at least needn’t) take any time for him to act on this reason.¹⁵ A father who notices that his daughter is in immediate danger acts instantly. These reasons, though they are excellent reasons to act, need not be

¹⁴ Of course, Fischer importantly does not require agents to be able to do otherwise in order to be morally responsible. So, this sentence is best interpreted as the claim that Schmidt’s voting mechanism wouldn’t have the capacity to properly weigh such information in that time.

¹⁵ Indeed, Fischer grants that the agoraphobe is reasons-responsive. And though Fischer argues that the agoraphobe is insufficiently reasons-responsive to warrant attributions of responsibility, this is because he is responsive to very few reasons, which is presumably not the case with Schmidt.

rigorously evaluated. They jive with these agents' normative orientations instantaneously. Surely we wouldn't want to claim these agents' mechanisms are not reasons-responsive. Indeed, those who act quickly in such situations are often praised, perhaps precisely because their decision mechanisms are highly reasons-responsive—they are able to quickly react to reasons without needing to ruminate. Though it's important to filter reasons through one's normative orientation, such filtration can take place almost instantaneously.

Given this, Schmidt may not require an extended period of time to analyze Block's inputs in order for his decision mechanism to be reasons-responsive. If Block's inputs were in-line with Schmidt's character, then it would be unnecessary for Schmidt to ponder them. For example, perhaps Block gave Schmidt several reasons to think his family would be in danger if he didn't vote for Hitler. Given that Schmidt cares deeply for his family, such reasons could act decisively, meshing quickly with his normative orientation. If Block's implanted reasons are in-line with Schmidt's character, then I see no reason to think that Schmidt's mechanism was any less reasons-responsive than it was in the case in which he votes for Hitler on his own. However, Long does not elaborate on what reasons Block provides Schmidt.

Of course, Fischer could have something else in mind entirely when he argues that Schmidt's voting mechanism isn't reasons-responsive. Fischer argues:

Anything that disrupts the typical relationship between an agent's existing normative orientation and a new element—sequestering the new element and rendering it immune from the causal interaction with the preexisting elements—cannot be the ordinary mechanism of human practical reasoning (Fischer 2012, 200).

While a short decision period could arguably stop reasons from interacting with one another (though as I argue above, it does not necessarily do so), there are other ways for this to occur. Perhaps the very nature of an input renders it causally inaccessible to preexisting elements. If the reasons Block implanted in Schmidt's brain were such that they *couldn't* be weighed against Schmidt's other reasons, then this really would seem to render his decision process non-reasons-responsive. While I'm not sure how one could implant such reasons in an agent, and these reasons certainly weren't featured in Long's case, Fischer's analysis has illuminated another way in which an agent can be rendered non-reasons-responsive.

Thus, in order to know if Schmidt's voting mechanism really was reasons-responsive in Long's manipulation cases, we must know more about these scenarios. If Block implanted reasons that meshed well with Schmidt's normative orientation and character, then it may very well be the case that Fischer is committed to the claim that Schmidt was responsible, despite having been manipulated. However, if Block implanted reasons that couldn't causally interact with any of Schmidt's other reasons, then Fischer can conclude that Schmidt was not morally responsible for his vote.

Interestingly, the under-description of these manipulation cases brings out an important aspect of accounts of guidance control. While Fischer is right to focus on the importance of how an agent responds to reasons, the reasons themselves play a crucially important part of reasons-responsiveness. Whether a reason is in-line with an agent's character or able to interact with the agent's other reasons affects an agent's ability to be morally responsible. And while I agree with both Long and Fischer that we can't be expected to control how we come to have the reasons that we have, I do think that the nature of these reasons affects our ability to respond to them, and thus affects the degree to which we are reasons-responsive.

IV. Conclusion: the Pilgrimage

In the introductory essay of this collection, Fischer offers a new guiding metaphor to understand agency and control. Typically, it's thought that exercising free agency is much like choosing between multiple paths; to be free is to be able to choose. Such a metaphor is incompatible with a view of control that does not require access to alternative possibilities. Even if our futures do resemble a garden of forking paths (Borges, 1941), for those who think regulative control isn't necessary for moral responsibility, this metaphor is highly uninformative. Thus, Fischer presents a new metaphor: the pilgrimage. For Fischer, to be a morally responsible agent is much like taking a pilgrimage. According to Fischer, "A pilgrimage is typically a preset path from one point to another; it does not—nor need not—involve alternate pathways ... Despite the fact that the pilgrimage route is laid out in advance, the pilgrims can achieve great personal growth and transformation" (Fischer 2012, 25). In other words, identifying with the choices one makes in life is what matters when it comes to agency and control, not whether we could have made other choices. Just like those who follow the set route of a pilgrimage, the value of our choices can't be

explained in terms of all the choices we forwent. Rather, it's the fact that they are *our* choices that gives them value.

I find the symbol of a pilgrimage illuminating. Indeed, many of the arguments Fischer lays out in *Deep Control* are captured nicely by this metaphor. In the first section, Fischer is able to demonstrate the irrelevance of regulative control by defending Frankfurt cases against numerous attacks. In doing so, Fischer is able to break away from the forking paths paradigm. In Chapter 10, Fischer rejects the idea that, in order to be responsible, agents must be somehow self-caused or the ultimate source of their behavior. He relies on Joel Feinberg's work on autonomy to explain why we need not operate with such a stringent notion of sourcehood and ultimacy. Feinberg writes: "Self-creation in the authentic person must be a process of self-*re*-creation, rationally accommodating new experiences and old policies to make greater coherence and flexibility. Self-creation is possible but not *ex nihilo*" (Feinberg 1986, 35). The idea that our agency arises from how we live coalesces nicely with the notion of a pilgrimage. The value of a pilgrimage is not that it was chosen over all other possible paths. Rather, the value lies in traveling it. Those who travel these paths do not begin them the same way they end them—their identities are transformed through their travels. The pilgrimage metaphor also grounds the notion of guidance control. To be moderately reasons-responsive is to travel the path adeptly, responsive to one's environment, and with one's destination in mind. Following a trail blindly and unquestioningly is to fail to be responsive to one's environment. And just as the unresponsive pilgrim gains little from his journey, the agent whose mechanisms fail to be moderately reasons-responsive cannot exercise guidance control. Furthermore, to exercise ownership over one's mechanisms is very much like identifying with one's pilgrimage. The pilgrim who is forced on her journey by nefarious forces will not only gain little from her travels, she won't be able to identify with them. Similarly, an agent who feels as though her thoughts and desires are not her own can neither identify with these mental states nor take ownership over them. In this way, by focusing our attention on how individuals travel a single path, as opposed to how they choose between them, we can see what really matters when it comes to acting freely. Fischer's collection features several extraordinary insights, many of which I have neglected in this review, but this shift from a forking paths perspective to one that focuses on a single journey is perhaps the most poignant. This shift allows us to look at the concept of freedom from a completely new angle, illuminating new features previously unnoticed.

V. Acknowledgements

I'd like to thank Michael McKenna for his guidance and his valuable remarks on earlier drafts of this essay. I am also grateful to John Martin Fischer, whose work on free will and moral responsibility has been extraordinarily influential in my thinking on these matters.

VI. Works Cited

- Ayer, A.J. 1954. "Freedom and Necessity," in A.J. Ayer, *Philosophical Essays*. New York: St. Martin's Press: 3-20.
- Campbell, J. 1997. "A Compatibilist Theory of Alternative Possibilities," *Philosophical Studies* 88: 319-330.
- Dennett, D. 1984. *Elbow Room: The Varieties of Free Will Worth Wanting*. Cambridge, Mass: MIT Press.
- Feinberg, J. 1986. *Harm to Self*. Volume Three: *The Moral Limits of the Law*. New York: Oxford University Press.
- Fischer, J.M. 2012. *Deep Control: Essays on Free Will and Value*, Oxford: Oxford University Press.
- _____. 2004. "Responsibility and Manipulation," *Journal of Ethics* 8: 145-177.
- _____. 1999. "Recent Work on Moral Responsibility," *Ethics* 110: 93-139.
- _____. 1982. "Responsibility and Control," *Journal of Philosophy* 79: 24-40.
- Frankfurt, H. 1969. "Alternate Possibilities and Moral Responsibility." *Journal of Philosophy* 66: 829-39.
- Ginet, C. 1996. "In Defense of the Principle of Alternative Possibilities: Why I Don't Find Frankfurt's Argument Convincing," *Philosophical Perspectives* 10: 403-417.
- Hobart, R.E. 1934. "Free Will as Involving Determinism and Inconceivable Without It," *Mind* 43: 1-27.
- Horgan, T. 1979. "'Could,' Possible Worlds, and Moral Responsibility," *Southern Journal of Philosophy* 17: 345-58.
- Hunt, D. 2000. "Moral Responsibility and Unavoidable Action," *Philosophical Studies*

- 97: 195-227.
- Kane, R. 1996. *The Significance of Free Will*. Oxford: Oxford University Press.
- Lehrer, K. 1976. "‘Can’ in Theory and Practice: A Possible Worlds Analysis," M. Brand and D. Walton (eds.) *Action Theory*: 241-70.
- Lewis, D. 1981. "Are We Free to Break the Laws?" *Theoria* 47: 113-21.
- Long, T. 2004. "Moderate Reasons-Responsiveness, Moral Responsibility, and Manipulation," in Joseph Campbell, Michael O’Rourke, and David Shier (eds.) *Freedom and Determinism*. Cambridge, MA: MIT Press, 151-172.
- McKenna, M. 2008. "A Hard-Line Reply to Pereboom’s Four-Case Manipulation Argument," *Phenomenological Research* 77: 142-159.
- _____. 2003. "Robustness, Control, and the Demand for Morally Significant Alternatives," in David Widerker and Michael McKenna (eds.) *Moral Responsibility and Alternative Possibilities*. Aldershot, UK: Ashgate Press.
- Mele, A. 2006. *Free Will and Luck*. New York: Oxford University Press.
- _____. 2000. "Reactive Attitudes, Reactivity, and Omissions," *Philosophy and Phenomenological Research* 61: 447-452.
- _____. and Robb, B. 1998. "Rescuing Frankfurt-Style Cases," *Philosophical Review* 107: 97-112.
- Nowell-Smith, P.H. 1948. "Free Will and Moral Responsibility," *Mind* 57: 45-61.
- Pereboom, D. 2001. *Living without Free Will*. Cambridge University Press.
- _____. 2000. "Alternate Possibilities and Causal Histories," in James Tomberlin (ed.) *Philosophical Perspectives Volume 20*. Astascadero: Ridgeview.
- Smart, J.J.C. 1961. "Free-Will, Praise, and Blame," *Mind* 70: 291-306.
- Speak, D. "The Impertinence of Frankfurt-style Argument," *The Philosophical Quarterly* 57: 76-95.
- Strawson, G. 2002. "The Bounds of Freedom," in Robert Kane (ed.) *The Oxford Handbook of Free Will*. Oxford: Oxford University Press.
- Van Inwagen, P. 1983. *An Essay on Free Will*. Oxford: Clarendon Press.
- Vihvelin, K. 2000. "Freedom, Foreknowledge, and the Principle of Alternate Possibilities." *Canadian Journal of Philosophy* 30: 1-24.
- Watson, G. 1998. "Some Worries about Semi-Compatibilism: Remarks on John Fischer’s

The Metaphysics of Free Will,” Journal of Social Philosophy 29: 135-143.

Widerker, D. 1995. “Libertarianism and Frankfurt’s Attack on the Principle of Alternative Possibilities.” *Philosophical Review* 104: 247-261.